arXiv:1703.05252v2 [math.CO] 26 Mar 2017

A Semi-Random Construction of Small Covering Arrays

Shagnik Das*† Tamás Mészáros*‡

March 28, 2017

Abstract

Given a set S of $v \geq 2$ symbols, and integers $k \geq t \geq 2$ and $N \geq 1$, an $N \times k$ array $A \in S^{N \times k}$ is an $(N; t, k, v)$-covering array if all sequences in S^t appear as rows in every $N \times t$ subarray of A. These arrays have a wide variety of applications, driving the search for small covering arrays. The covering array number, $\mathrm{CAN}(t, k, v)$, is the smallest N for which an $(N; t, k, v)$-covering array exists.

In this paper, we combine probabilistic and linear algebraic constructions to improve the upper bounds on $\mathrm{CAN}(t, k, v)$ by a factor of $\ln v$, showing that for prime powers v, $\mathrm{CAN}(t, k, v) \leq (1 + o(1))\left((t-1)v^t/(2\log_2 v - \log_2(v+1))\right)\log_2 k$, which also offers improvements for large v that are not prime powers. Our main tool, which may be of independent interest, is a construction of an array with v^t rows that covers the maximum possible number of subsets of size t.

1 Introduction

In the last few decades, a great deal of research has been devoted to the study of orthogonal and covering arrays, an important class of combinatorial designs. This research is motivated by numerous applications, in particular to computer science and the design of experiments, and one of the major open problems in this area is to determine how small these arrays can be. In this paper we improve the best-known general bounds for this problem. We first provide a brief introduction to the subject, surveying some relevant results from the literature, before presenting our new results.

1.1 Background and previous results

Let A be an $N \times k$ array, whose entries come from some set S of v symbols; that is, $A \in S^{N \times k}$. In the context of experimental design, N represents the number of trials to be carried out, k denotes the number of factors to be tested, and S is the set of levels these factors can take.

*Freie Universität Berlin, Institut für Mathematik, Arnimallee 3, 14195 Berlin, Germany.
†E-mail: shagnik@mi.fu-berlin.de
‡E-mail: tamas.meszaros@fu-berlin.de. Position funded by the DRS POINT Fellowship Program.

The objective is to determine how subsets of the factors interact with one another. To that end, given a set Q of t column indices, we denote by A_Q the $N \times t$ subarray obtained by restricting A to the columns in Q.

Definition 1.1 (Orthogonal arrays)**.** *Given a set S of v symbols, k columns, an* index $\lambda \in \mathbb{N}$ *and a* strength $t \in \mathbb{N}$*, let $N = \lambda v^t$. An $N \times k$ array $A \in S^{N \times k}$ is an* $(N; t, k, v)$-orthogonal array *if for every subset Q of t columns, every sequence in S^t appears exactly λ times as a row of the subarray A_Q.*

An orthogonal array is therefore a very regular structure, behaving uniformly with respect to every subset of t columns, giving rise to an important application in theoretical computer science. Many randomised algorithms use some large number k of independent random variables, each uniformly distributed over a set S of size v, thus using the exponentially large probability space S^k. However, quite often one only requires the weaker condition that the random variables be t-wise independent, for some small t. Given an $(N; t, k, v)$-orthogonal array, a uniform distribution on the N rows of this array provides a probability space with the desired independence, and, if N is small, this allows for brute-force derandomisation of the algorithm.

The interest, then, is in determining how few rows an orthogonal array with a given strength and number of columns can have. In one of the early papers on the subject, Plackett and Burman [22] provided sharp bounds for orthogonal arrays of strength two, showing how the number of rows must grow with the number of columns.

Theorem 1.2 (Plackett–Burman [22], 1946)**.** *If an $(N; 2, k, v)$-orthogonal array exists, we must have*

$$k \leq \left\lfloor \frac{N-1}{v-1} \right\rfloor.$$

Following this initial focus on orthogonal arrays of strength two, Rao [23, 24] generalised the notion to arrays of strength t, giving rise to the modern study of orthogonal arrays. For an account of the last half-century's developments in the field, the reader is referred to the book of Hedayat, Sloane and Stufken [13].

However, the high level of regularity required of an orthogonal array places severe restrictions on the possible values of the parameters, and hence the fundamental question asks for which parameters an $(N; t, k, v)$-orthogonal array exists. For many applications, one is willing to make do with a smaller, yet less regular, construction, giving rise to the relaxation of orthogonal arrays to covering arrays. Here, one only requires that all sequences in S^t appear at least once in every $N \times t$ subarray, dropping the condition that they appear equally often. The primary question is now an extremal one — how small can an array satisfying this weaker condition be?

Definition 1.3 (Covering arrays)**.** *Given a set S of v symbols, N rows, k columns and a* strength $t \in \mathbb{N}$*, let $A \in S^{N \times k}$ be an $N \times k$ array. We say that A* covers *a subset Q of t columns if every sequence in S^t appears at least once as a row of the subarray A_Q.*

The array A is an $(N; t, k, v)$-covering array *if it covers every t-subset of the k columns, and the* covering array number $\mathrm{CAN}(t, k, v)$ *is the minimum number N of rows for which an $(N; t, k, v)$-covering array exists.*

In the binary case $v = 2$, S may be taken to be the set $\{0, 1\}$, and an $(N; t, k, 2)$-covering array can then be interpreted as the incidence matrix A of a family $\mathcal{F}$ of N subsets of the ground set $[k]$. In this setting, covering is often referred to as *shattering*, as a t-subset $Q \subseteq [k]$ is shattered whenever every one of its 2^t subsets appears as an intersection of Q with some set F in the family $\mathcal{F}$. If the array is an $(N; t, k, 2)$-covering array, the corresponding family is said to be (k, t)-universal.

The study of such families dates back to the early 1970's, when one of the few exact results in this field was obtained by Rényi [25] and Katona [14], and independently by Kleitman and Spencer [15]. In our terminology, they showed

$$\mathrm{CAN}(2, k, 2) = \min\left\{N : k \leq \binom{N-1}{\lfloor \frac{N}{2} \rfloor - 1}\right\} = \log_2 k + \left(\frac{1}{2} + o(1)\right) \log_2 \log_2 k.$$

For larger t, Kleitman and Spencer [15] showed there are constants $c_1, c_2 > 0$ such that

$$c_1 2^t \log_2 k \leq \mathrm{CAN}(t, k, 2) \leq c_2 t 2^t \log_2 k.$$

These results have since been extended to covering arrays over larger sets of symbols. In the strength-two case, Gargano, Körner and Vaccaro [10] established the asymptotic result

$$\mathrm{CAN}(2, k, v) = \left(\frac{1}{2} + o(1)\right) v \log_2 k.$$

For higher strengths, there is a considerable gap between the best-known upper and lower bounds on $\mathrm{CAN}(t, k, v)$. A general upper bound was given by Godbole, Skipper and Sunley [11], who studied when the uniformly random array $A \in S^{N \times k}$ is an $(N; t, k, v)$-covering array.

Theorem 1.4 (Godbole–Skipper–Sunley [11], 1996)**.** *For fixed integers $t, v \geq 2$, as $k \to \infty$,*

$$\mathrm{CAN}(t, k, v) \leq (1 + o(1)) \frac{(t-1) \log_2 k}{\log_2 \frac{v^t}{v^t - 1}}.$$

One can prove a lower bound by induction on t, reducing the problem to the $t = 2$ case. This shows $\mathrm{CAN}(t, k, v) \geq \left(\frac{1}{2} + o(1)\right) v^{t-1} \log_2(k - t + 2)$. In particular, when t and v are fixed, $\mathrm{CAN}(t, k, v) = \Theta_{t,v}(\log_2 k)$, but the correct dependence on t and v is unknown. For convenience, we often consider the behaviour of the function in the limit as k goes to infinity.

Definition 1.5. *Given integers $t, v \geq 2$, we define*

$$d(t, v) = \limsup_{k \to \infty} \frac{\mathrm{CAN}(t, k, v)}{\log_2 k}.$$

The result of Gargano, Körner and Vaccaro [10] thus shows that $d(2, v) = \frac{v}{2}$, while Theorem 1.4 can be restated as saying $d(t, v) \leq \frac{t-1}{\log_2 \frac{v^t}{v^t-1}}$. Note that in the limit as $v^t \to \infty$, this bound is asymptotically $(t-1)v^t \ln 2$.

There have since been some improvements to the lower-order terms of this upper bound. Francetić and Stevens [7] showed that for all $t, v \geq 2$,

$$d(t, v) \leq \frac{v(t-1)}{\log_2 \frac{v^{t-1}}{v^{t-1}-1}}.$$

Sarkar and Colbourn [27] gave an alternative proof of this improved bound, and further reduced the bound whenever v is a prime power. In this case, they proved

$$d(t, v) \leq \frac{v(v-1)(t-1)}{\log \frac{v^{t-1}}{v^{t-1}-v+1}}.$$

While these bounds give the same asymptotics as Theorem 1.4 as v^t tends to infinity, they are significantly better for small values of v and t. For further details of the best-known constructions for particular values of v or t, we refer the reader to the excellent surveys of Lawrence, Kacker, Lei, Kuhn and Forbes [16] in the binary setting and of Colbourn [5] in the general setting.

1.2 Our results

We begin by studying what is, in some sense, an inverse problem. Rather than seeking the smallest array that covers all t-sets of columns, we fix the size of the array and try to maximise the number of covered t-sets. This gives rise to the following extremal function.

Definition 1.6 (Maximum coverage function)**.** *Suppose we have a set S of v symbols, N rows, k columns and a strength $t \in \mathbb{N}$. For an $N \times k$ array $A \in S^{N \times k}$, let* $\mathrm{Cov}(A)$ *denote the collection of all subsets of columns that are covered by A, and define* $\mathrm{cov}_t(A)$ *to be the number of sets of size t in* $\mathrm{Cov}(A)$*.*

We define the maximum coverage function, $\mathrm{cov}_{\max}(N; t, k, v)$*, to be the maximum number of t-subsets that can be covered by such an array. That is,*

$$\mathrm{cov}_{\max}(N; t, k, v) = \max\left\{\mathrm{cov}_t(A) : A \in S^{N \times k}\right\}.$$

We note that similar notions have appeared previously in the literature. Hartman and Raskin [12] and Maximoff, Trela, Kuhn and Kacker [17] suggested comparable lines of study, with a focus on developing heuristics for building small arrays that cover many sets. On the other hand, our interest is in proving general bounds on the function $\mathrm{cov}_{\max}(N; t, k, v)$. More recently, Sarkar, Colbourn, De Bonis and Vaccaro [28] proved bounds on the sizes of almost-covering arrays, which are arrays that cover almost all t-sets. This is more closely related to our investigation, and we shall discuss their results in our concluding remarks.

Note that in order for an array $A \in S^{N\times k}$ to cover even a single t-set Q of columns, we must have at least v^t rows, as each of the v^t sequences in S^t must appear as rows in A_Q. This trivially gives $\mathrm{cov}_{\max}(N;t,k,v) = 0$ for all $N \leq v^t - 1$. The first problem of interest is thus to determine $\mathrm{cov}_{\max}(v^t;t,k,v)$, and this is the case on which we focus.

Given that the upper bound of Theorem 1.4 comes from the random array, this is a natural candidate to consider for our problem as well. Let A_{rand,v^t} be a uniformly random array chosen from $S^{v^t\times k}$. Observe that for any subset Q of t columns, the v^t rows of $(A_{\mathrm{rand},v^t})_Q$ are independent and uniformly distributed over S^t. The probability that these rows are all distinct, and hence that Q is covered, is thus

$$\mathbb{P}\left(Q \in \mathrm{Cov}(A_{\mathrm{rand},v^t})\right) = \frac{(v^t)!}{(v^t)^{v^t}}.$$

As v^t tends to infinity, this is $e^{-(1-o(1))v^t}$, and so in expectation A_{rand,v^t} only covers an exponentially small fraction of all t-sets of columns.

A moment's thought reveals that this is far from optimal. Indeed, consider the following block construction, where for simplicity we suppose that k is divisible by t. Partition the k columns into t equal-sized subsets $B_1, B_2, \ldots, B_t$, and build an array A_{block} whose rows are all sequences in S^k that are constant on the sets B_i, $1 \leq i \leq t$. We clearly have exactly v^t rows, and a set Q of columns is covered if and only if $|Q \cap B_i| \leq 1$ for all i. Therefore, when k and t are suitably large, we have

$$\mathrm{cov}_t(A_{\mathrm{block}}) = \binom{k}{t}^t \sim \sqrt{2\pi t}\cdot e^{-t}\binom{k}{t}.$$

This already gives a significantly better lower bound on $\mathrm{cov}_{\max}(v^t;t,k,v)$ than the random array A_{rand,v^t}, but, as it turns out, is still far from the truth. In our primary result below we give an upper bound on $\mathrm{cov}_{\max}(v^t;t,k,v)$ and, whenever v is a prime power, demonstrate its tightness by means of an explicit construction.

Theorem 1.7. *Given integers $t, v \geq 2$, define*

$$c_{t,v} = \prod_{i=0}^{t-1} \frac{v^t - v^i}{v^t - 1}.$$

Then, for any $k \geq t$, we have the bound

$$\mathrm{cov}_{\max}(v^t;t,k,v) \leq c_{t,v}\frac{k^t}{t!}.$$

If v is a prime power, and $\frac{v^t-1}{v-1}$ divides k, we have equality; that is, $\mathrm{cov}_{\max}(v^t;t,k,v) = c_{t,v}\frac{k^t}{t!}$. For prime power v and all $k \geq t$, we have

$$\mathrm{cov}_{\max}(v^t;t,k,v) \geq c_{t,v}\binom{k}{t}.$$

Observe that

$$c_{t,v} = \prod_{i=0}^{t-1} \frac{v^t - v^i}{v^t - 1} \geq \prod_{i=1}^{t-1} \left(1 - v^{-i}\right) \geq \left(1 - v^{-1}\right) \left(1 - \sum_{i=2}^{\infty} v^{-i}\right) = 1 - \frac{v+1}{v^2}.$$

This shows, perhaps surprisingly, that as soon as we have a large enough array to cover a single t-set of columns, one can already cover a large proportion of all such t-sets. We can then use these very efficient arrays to build small covering arrays.

Corollary 1.8. *For integers $k \geq t \geq 2$ and a prime power $v \geq 2$, we have*

$$\mathrm{CAN}(t,k,v) \leq \left\lceil \frac{(t-1)\log_2 k + \log_2(et)}{\log_2 \frac{1}{1-c_{t,v}}} \right\rceil v^t.$$

In particular, this implies that whenever v is a prime power,

$$d(t,v) \leq \frac{(t-1)v^t}{\log_2 \frac{1}{1-c_{t,v}}} \leq \frac{(t-1)v^t}{2\log_2 v - \log_2(v+1)}.$$

Moreover, there is some absolute constant $v_0 \in \mathbb{N}$ such that for all integers $v \geq v_0$, we have

$$d(t,v) \leq \frac{(t-1)e^{tv^{-0.474}}v^t}{2\log_2 v - \log_2(v+1)}.$$

This represents the first asymptotic improvement on the bound in Theorem 1.4. Recall that, as v^t grows, the bounds from Theorem 1.4 and the subsequent improvements are all asymptotically $(t-1)v^t \ln 2$, and hence our bound is smaller by a factor of approximately $\ln v$ (provided that either v is large and $t = o(v^{0.474})$, or v is a prime power). More precise calculations suggest that for prime power v, the bounds in Corollary 1.8 improve the existing bounds whenever $t \geq 3$ and $v \geq 4$.

1.3 Organisation and notation

The remainder of this paper is organised as follows. In Section 2 we prove the upper bound of Theorem 1.7, and in Section 3 we establish the lower bound by providing an optimal construction. Section 4 is devoted to the construction of small covering arrays and the proof of Corollary 1.8. We then provide some concluding remarks and open problems in Section 5.

We use standard combinatorial notation throughout this paper. In particular, $[k]$ denotes the first k positive integers, $\{1, 2, \ldots, k\}$. Given a set X and an integer t, $\binom{X}{t}$ is the collection of all t-subsets of X. As defined previously, for an array A and a subset Q of its columns, A_Q denotes the subarray of A containing only the columns in Q. We denote by $\mathrm{Cov}(A)$ the collection of all subsets of columns that are covered by A, and by $\mathrm{cov}_t(A)$ the number of sets of size t in $\mathrm{Cov}(A)$. Finally, we use $\log_2$ for the binary logarithm, and ln for the natural logarithm.

2 A general upper bound

We shall start our proof of Theorem 1.7 by proving the general upper bound. In order to do this, we shall show that the number of t-sets that can be covered by an array of size $v^t \times k$ is bounded by the Lagrangian of an auxiliary hypergraph, which we shall then bound. First we present some useful preliminaries concerning Lagrangians in general.

2.1 Lagrangians of t-uniform hypergraphs

Lagrangians, first introduced by Motzkin and Straus [20] to give a proof of Turán's theorem, have proven to be very useful in the study of extremal combinatorics. Roughly speaking, the Lagrangian of a hypergraph determines the maximum possible density of a blow-up of the hypergraph. We now define the Lagrangian more precisely.

Definition 2.1. *Let $\mathcal{H} \subseteq \binom{U}{t}$ be a t-uniform hypergraph on a finite set U of vertices. We say that a function $x : U \to \mathbb{R}$ is a legal weighting if $x(u) \geq 0$ for every $u \in U$ and $\sum_{u \in U} x(u) = 1$.*

The weight polynomial of $\mathcal{H}$ evaluated at this weighting is given by

$$w(x, \mathcal{H}) = \sum_{e \in \mathcal{H}} \prod_{u \in e} x(u).$$

The Lagrangian of $\mathcal{H}$ is defined to be $\lambda(\mathcal{H}) = \max w(x, \mathcal{H})$, where the maximum is taken over all legal weightings x of $\mathcal{H}$.

We call a legal weighting x optimal if $w(x, \mathcal{H}) = \lambda(\mathcal{H})$. The following lemma, proven by Frankl and Rödl [8], gives some information about the minimal supports of optimal weightings.

Lemma 2.2. *Let $\mathcal{H}$ be a t-uniform hypergraph on the vertex set U, and suppose x is an optimal weighting where the number of vertices with non-zero weight is minimal. If $u, w \in U$ are vertices of non-zero weight, then there is an edge in $\mathcal{H}$ containing both u and w.*

Proof. Suppose for contradiction that vertices u and w have positive weight, but there is no edge of $\mathcal{H}$ containing both of them. We define a parametrised function $x_\varepsilon : U \to \mathbb{R}$, where

$$x_\varepsilon(s) = \begin{cases} x(u) + \varepsilon & \text{if } s = u, \\ x(w) - \varepsilon & \text{if } s = w, \\ x(s) & \text{otherwise.} \end{cases}$$

Observe that x_ε is a legal weighting whenever $\varepsilon \in [-x(u), x(w)]$, and that at the boundaries either $x_\varepsilon(u)$ or $x_\varepsilon(w)$ becomes zero. Moreover, since there is no edge containing both u and w, $w(x_\varepsilon, \mathcal{H})$ is linear in ε, and hence is maximised by some $\varepsilon^* \in \{-x(u), x(w)\}$. However, this gives a contradiction, as x_{ε^*} is then an optimal weighting with fewer non-zero weights. Thus u and w must appear together in some edge of $\mathcal{H}$. □

The next lemma shows that we can bound the Lagrangian $\lambda(\mathcal{H})$ of a t-uniform hypergraph $\mathcal{H}$ in terms of the Lagrangians of smaller $(t-1)$-uniform hypergraphs. Given a vertex $u \in U$, the *link hypergraph* $\mathcal{H}(u)$ is a $(t-1)$-uniform hypergraph on the vertex set $U \setminus \{u\}$, with edges $e' \in \mathcal{H}(u)$ whenever $e' \cup \{u\} \in \mathcal{H}$.

Lemma 2.3. *Given $t \geq 2$, a t-uniform hypergraph $\mathcal{H}$ on the vertex set U and a legal weighting x, we have*

$$w(x, \mathcal{H}) \leq \frac{1}{t} \sum_{u \in U} x(u)(1 - x(u))^{t-1} \lambda(\mathcal{H}(u)).$$

Proof. Note that if $x(u) = 1$ for any $u \in U$, then $w(x, \mathcal{H}) = 0$, and the inequality trivially holds. Hence we may assume $x(u) < 1$ for all $u \in U$. Since every edge of $\mathcal{H}$ has exactly t vertices, double-counting gives

$$w(x, \mathcal{H}) = \sum_{e \in \mathcal{H}} \prod_{u \in e} x(u) = \frac{1}{t} \sum_{u \in U} \sum_{u \in e \in \mathcal{H}} \prod_{w \in e} x(w) = \frac{1}{t} \sum_{u \in U} x(u) \sum_{e' \in \mathcal{H}(u)} \prod_{w \in e'} x(w).$$

As x is a legal weighting, for every vertex u, x must distribute a total weight of $1 - x(u)$ on the vertices in $U \setminus \{u\}$. Thus we can rescale the weights to obtain a legal weighting for the link hypergraph $\mathcal{H}(u)$ by defining $x_u(w) = \frac{x(w)}{1-x(u)}$ for all vertices $w \in U \setminus \{u\}$. We then have

$$\begin{aligned} w(x, \mathcal{H}) &= \frac{1}{t} \sum_{u \in U} x(u) \sum_{e' \in \mathcal{H}(u)} \prod_{w \in e'} x(w) = \frac{1}{t} \sum_{u \in U} x(u)(1 - x(u))^{t-1} \sum_{e' \in \mathcal{H}(u)} \prod_{w \in e'} x_u(w) \\ &= \frac{1}{t} \sum_{u \in U} x(u)(1 - x(u))^{t-1} w(x_u, \mathcal{H}(u)) \leq \frac{1}{t} \sum_{u \in U} x(u)(1 - x(u))^{t-1} \lambda(\mathcal{H}(u)), \end{aligned}$$

where the inequality follows from the definition of the Lagrangian. □

The final lemma of this subsection concerns the Lagrangian of the v-fold tensor product of a hypergraph with itself. Given t-uniform hypergraphs, $\mathcal{H}_1, \ldots, \mathcal{H}_v$ on vertex sets $U_1, \ldots, U_v$ respectively, their tensor product, denoted $\otimes_{\ell=1}^{v} \mathcal{H}_\ell$, is a t-uniform hypergraph on the vertex set $\prod_{\ell=1}^{v} U_\ell$ with edges

$$\otimes_{\ell=1}^{v} \mathcal{H}_\ell = \left\{ \{(u_{1,1}, \ldots, u_{v,1}), \ldots, (u_{1,t}, \ldots, u_{v,t})\} : \forall \ell \in [v],\ \{u_{\ell,1}, \ldots, u_{\ell,t}\} \in \mathcal{H}_\ell \right\}.$$

Note that every t-tuple of edges $(e_1, \ldots, e_v) \in \prod_{\ell=1}^{v} \mathcal{H}_\ell$ gives rise to $(t!)^{v-1}$ edges in $\otimes_{\ell=1}^{v} \mathcal{H}_\ell$, as the vertices can be combined in every possible order. We write $\mathcal{H}^{\otimes v}$ for the hypergraph obtained by taking a tensor product of v copies of a hypergraph $\mathcal{H}$.

Lemma 2.4. *For every t-uniform hypergraph $\mathcal{H}$,*

$$\lambda(\mathcal{H}^{\otimes v}) = \lambda(\mathcal{H}).$$

Proof. Let U be the vertex set of $\mathcal{H}$. First consider an optimal weighting x for $\mathcal{H}$, and define a legal weighting $\widehat{x}$ on U^v as follows:

$$\widehat{x}((u_1,u_2,\ldots,u_v)) = \begin{cases} x(u_1) & \text{if } u_1=u_2=\ldots=u_v,\\ 0 & \text{otherwise.}\end{cases}$$

Now for every edge $\{u_1,\ldots,u_t\}$ in $\mathcal{H}$, the edge $\{(u_1,\ldots,u_1),\ldots,(u_t,\ldots,u_t)\}$ appears in $\mathcal{H}^{\otimes v}$ with the same weight, and so

$$\lambda(\mathcal{H}^{\otimes v}) \geq w(\widehat{x},\mathcal{H}^{\otimes v}) \geq w(x,\mathcal{H}) = \lambda(\mathcal{H}).$$

For the reverse inequality, let $\widehat{x}$ be an optimal weighting for $\mathcal{H}^{\otimes v}$, and define the legal weighting x on U by $x(u)=\sum_{u_2,\ldots,u_v\in U}\widehat{x}((u,u_2,\ldots,u_v))$. We then have

$$\begin{aligned}\lambda(\mathcal{H}) \geq w(x,\mathcal{H}) &= \sum_{\{u_1,\ldots,u_t\}\in\mathcal{H}}\prod_{i=1}^{t}x(u_i) = \sum_{\substack{\{u_1,\ldots,u_t\}\in\mathcal{H}\\ u_{j,i}\in U,\ 2\leq j\leq v,\ i\in[t]}}\prod_{i=1}^{t}\widehat{x}((u_i,u_{2,i},\ldots,u_{v,i}))\\ &\geq \sum_{\substack{\{u_1,\ldots,u_t\}\in\mathcal{H}\\ \{u_{j,1},\ldots,u_{j,t}\}\in\mathcal{H},\ 2\leq j\leq v}}\prod_{i=1}^{t}\widehat{x}((u_i,u_{2,i},\ldots,u_{v,i}))\\ &= \sum_{\{(u_1,u_{2,1},\ldots,u_{v,1}),\ldots,(u_t,u_{2,t},\ldots,u_{v,t})\}\in\mathcal{H}^{\otimes v}}\prod_{i=1}^{t}\widehat{x}((u_i,u_{2,i},\ldots,u_{v,i})) = w(\widehat{x},\mathcal{H}^{\otimes v}) = \lambda(\mathcal{H}^{\otimes v}),\end{aligned}$$

and thus $\lambda(\mathcal{H})=\lambda(\mathcal{H}^{\otimes v})$. □

2.2 The upper bound

With these preliminaries in place, we may proceed with the proof of the upper bound in Theorem 1.7. As mentioned earlier, we shall bound the number of t-sets that can be covered by the Lagrangian of an auxiliary hypergraph $\mathcal{H}_{t,v}$, which we now introduce.

The vertex set of $\mathcal{H}_{t,v}$ is $[v]^{v^t}$ and t vectors $\vec{y}_1,\ldots,\vec{y}_t\in[v]^{v^t}$ form an edge in $\mathcal{H}_{t,v}$ whenever all v^t vectors in $[v]^t$ appear as rows of the $v^t\times t$ matrix whose columns are $\vec{y}_1,\ldots,\vec{y}_t$. Note that this condition implies the vectors $\vec{y}_i$ are pairwise-distinct, so this is indeed a t-uniform hypergraph.

For a $v^t\times k$ array A, whose entries we shall assume to belong to $[v]$, and $\vec{y}\in[v]^{v^t}$, define $B_{\vec{y}}=\{a\in[k]: A_{\{a\}}=\vec{y}\}$ and note that the sets $B_{\vec{y}}$ partition the set $[k]$ of columns of A. Observe that a t-set $Q\subseteq[k]$ can only be covered by A if all elements of Q belong to different parts $B_{\vec{y}}$, as otherwise two of the columns in A_Q will be identical. However, not all such t-sets are covered. A t-set $Q\subseteq[k]$ is covered by A if and only if the columns of A_Q form an edge of $\mathcal{H}_{t,v}$.

We can hence count the number of covered t-sets, finding

$$\mathrm{cov}_t(A) = \sum_{e\in\mathcal{H}_{t,v}}\prod_{\vec{y}\in e}|B_{\vec{y}}|\,.$$

Since the sets $\left\{B_{\vec{y}} : \vec{y} \in [v]^{v^t}\right\}$ partition $[k]$, the function $x : [v]^{v^t} \to \mathbb{R}$ given by $x(\vec{y}) = \frac{1}{k}\left|B_{\vec{y}}\right|$ is a legal weighting of the vertices of $\mathcal{H}_{t,v}$. Hence

$$\mathrm{cov}_t(A) = \sum_{e \in \mathcal{H}_{t,v}} \prod_{\vec{y} \in e} |B_{\vec{y}}| = k^t \sum_{e \in \mathcal{H}_{t,v}} \prod_{\vec{y} \in e} x(\vec{y}) = k^t w(x, \mathcal{H}_{t,v}) \leq k^t \lambda(\mathcal{H}_{t,v}).$$

Now note that the hypergraph $\mathcal{H}_{t,v}$ is in fact independent of the array A (which only determines the weighting of the vertices), and hence $k^t\lambda(\mathcal{H}_{t,v})$ bounds the number of t-sets that can be covered by any array of size $v^t \times k$. The following proposition therefore gives the desired upper bound for Theorem 1.7.

Proposition 2.5. *For all $t \geq 1$,*

$$\lambda(\mathcal{H}_{t,v}) \leq \frac{c_{t,v}}{t!} = \frac{1}{t!}\prod_{i=0}^{t-1}\frac{v^t - v^i}{v^t - 1}.$$

We first establish a few simple lemmas that we shall use when proving Proposition 2.5. The first shows that there are optimal weightings of $\mathcal{H}_{t,v}$ that are supported on a relatively small number of vertices.

Lemma 2.6. *There is an optimal weighting x of $\mathcal{H}_{t,v}$ for which*

$$|\mathrm{supp}(x)| = \left|\left\{\vec{y} \in [v]^{v^t} : x(\vec{y}) \neq 0\right\}\right| \leq \frac{v^t - 1}{v - 1}.$$

Proof. Let x be an optimal weighting of $\mathcal{H}_{t,v}$ minimising the number of vectors with non-zero weight. By Lemma 2.2, if $\vec{y}$ and $\vec{z}$ are vectors with non-zero weight, then there must be an edge $e \in \mathcal{H}_{t,v}$ containing both $\vec{y}$ and $\vec{z}$. Since every vector in $[v]^t$ appears as a row in the matrix whose columns are the vectors in e, it follows that for each choice of $a, b \in [v]$, there are exactly v^{t-2} coordinates i where $y_i = a$ and $z_i = b$. In particular, the vectors of non-zero weight form the columns of an $(v^t; 2, k, v)$-orthogonal array, and so by Theorem 1.2 their number is bounded from above by $\frac{v^t-1}{v-1}$. □

The next lemma describes the link hypergraphs of $\mathcal{H}_{t,v}$.

Lemma 2.7. *For any vertex $\vec{y} \in [v]^{v^t}$ in $\mathcal{H}_{t,v}$ of positive degree, $\mathcal{H}_{t,v}(\vec{y}) \cong \mathcal{H}_{t-1,v}^{\otimes v}$.*

Proof. Recall that we have an edge $\{\vec{y}_1, \ldots, \vec{y}_{t-1}, \vec{y}\} \in \mathcal{H}_{t,v}$ if and only if the $v^t \times t$ matrix M formed with these column vectors has all vectors in $[v]^t$ as row vectors. In particular, $\vec{y}$ must have v^{t-1} entries equal to a for every $a \in [v]$. Given $a \in [v]$, denote by $M^{(a)}$ the $v^{t-1} \times (t-1)$ matrix formed by taking those rows of M that end with a, and then deleting the last (all-a) column. Since the rows of M contain every vector of $[v]^t$ ending in a, it follows that the rows of $M^{(a)}$ consist of all vectors in $[v]^{t-1}$. In particular, the columns of $M^{(a)}$ form an edge in $\mathcal{H}_{t-1,v}$.

In other words, after a possible reordering of the rows of M, for $1 \leq i \leq t-1$ the vector $\vec{y}_i = (\vec{w}_{1,i}, \vec{w}_{2,i}, \ldots, \vec{w}_{v,i})^T$ should be the concatenation of vectors $\vec{w}_{j,i} \in [v]^{v^{t-1}}$, $j \in [v]$, such that for $1 \leq j \leq v$ we have $\{\vec{w}_{j,1}, \ldots, \vec{w}_{j,t-1}\} \in \mathcal{H}_{t-1,v}$. This correspondence between the edges $\{\vec{y}_1, \ldots, \vec{y}_{t-1}\} \in \mathcal{H}_{t,v}(\vec{y})$ and $\{(\vec{w}_{1,1}, \ldots, \vec{w}_{v,1}), \ldots, (\vec{w}_{1,t-1}, \ldots, \vec{w}_{v,t-1})\} \in \mathcal{H}_{t-1,v}^{\otimes v}$ gives the desired isomorphism between $\mathcal{H}_{t,v}(\vec{y})$ and $\mathcal{H}_{t-1,v}^{\otimes v}$. □

The final lemma solves an optimisation problem that shall appear in our proof of Proposition 2.5.

Lemma 2.8. *For $K \geq t \geq 2$, let $f(x_1, \ldots, x_K) = \sum_{i=1}^{K} x_i(1-x_i)^{t-1}$. The maximum of f, subject to $x_i \geq 0$ for every $1 \leq i \leq K$ and $\sum_i x_i = 1$, is $\left(1 - \frac{1}{K}\right)^{t-1}$.*

Proof. Note that f is a continuous function and the constraints define a compact set, so the maximum is well defined. By taking $x_i = \frac{1}{K}$ for all i, we find $f\left(\frac{1}{K}, \ldots, \frac{1}{K}\right) = \left(1 - \frac{1}{K}\right)^{t-1}$, and so we only need to prove the upper bound.

Set $g(x) = x(1-x)^{t-1}$, and observe that

$$g'(x) = (1-tx)(1-x)^{t-2} \quad \text{and} \quad g''(x) = (tx-2)(t-1)(1-x)^{t-3}.$$

Hence $g(x)$ is monotone increasing on $[0, \frac{1}{t}]$, monotone decreasing on $[\frac{1}{t}, 1]$, concave on $[0, \frac{2}{t}]$ and convex on $[\frac{2}{t}, 1]$.

Now let $(x_1, \ldots, x_K)$ maximise f, and suppose there was some index i_0 such that $x_{i_0} > \frac{1}{t}$. Since $K \geq t$ and $\sum_i x_i = 1$, the average of the weights x_i is at most $\frac{1}{t}$, and hence there must be some index j_0 such that $x_{j_0} < \frac{1}{t}$. Put $\varepsilon = \min\left\{\frac{1}{t} - x_{j_0}, x_{i_0} - \frac{1}{t}\right\} > 0$. If we replace x_{j_0} by $x_{j_0} + \varepsilon$ and x_{i_0} by $x_{i_0} - \varepsilon$, the monotonicity properties of g imply the value of f would increase, contradicting the fact that $(x_1, \ldots, x_K)$ maximises f.

Hence $x_i \in \left[0, \frac{1}{t}\right]$ for $1 \leq i \leq K$. As g is concave on this interval, Jensen's inequality gives

$$f(x_1, \ldots, x_K) = \sum_{i=1}^{K} x_i(1-x_i)^{t-1} = \sum_{i=1}^{K} g(x_i) \leq Kg\left(\frac{1}{K}\sum_{i=1}^{K} x_i\right) = \left(1 - \frac{1}{K}\right)^{t-1}. \qquad \square$$

We are now in position to bound the Lagrangian of the hypergraph $\mathcal{H}_{t,v}$, thereby completing the proof of the upper bound from Theorem 1.7.

Proof of Proposition 2.5. For fixed v, we shall prove $\lambda(\mathcal{H}_{t,v}) \leq \frac{c_{t,v}}{t!}$ by induction on t.

For the base case $t = 1$, the hypergraph $\mathcal{H}_{v,1}$ is very simple. We have v^v vertices corresponding to the vectors in $[v]^v$. The edges of $\mathcal{H}_{v,1}$ are the singletons corresponding to vectors containing every $a \in [v]$. We thus have $v!$ edges, and the weight polynomial is simply the sum of the weights of the corresponding $v!$ vertices, whose maximum value is trivially at most 1, which is equal to $\frac{c_{1,v}}{1!}$.

For the induction step, suppose $\lambda(\mathcal{H}_{t-1,v}) \leq \frac{c_{t-1,v}}{(t-1)!}$ and consider $\mathcal{H}_{t,v}$. Let x be an optimal weighting of $\mathcal{H}_{t,v}$ with minimal support. Suppose $\vec{y}_1, \ldots, \vec{y}_K$ are the vertices of $\mathcal{H}_{t,v}$ with non-zero weight, and let $x_1, \ldots, x_K$ represent their respective weights. By Lemma 2.6, $K \leq \frac{v^t - 1}{v-1}$.

Using Lemma 2.3, we find

$$\lambda(\mathcal{H}_{t,v}) = w(x, \mathcal{H}_{t,v}) \leq \frac{1}{t}\sum_{\vec{y} \in [v]^{v^t}} x(\vec{y})(1 - x(\vec{y}))^{t-1}\lambda(\mathcal{H}_{t,v}(\vec{y})) = \frac{1}{t}\sum_{i=1}^{K} x_i(1-x_i)^{t-1}\lambda(\mathcal{H}_{t,v}(\vec{y}_i)).$$

By Lemma 2.7, $\mathcal{H}_{t,v}(\vec{y}_i) \cong \mathcal{H}_{t-1,v}^{\otimes v}$ for all i, and hence Lemma 2.4 and the induction hypothesis give $\lambda(\mathcal{H}_{t,v}(\vec{y}_i)) = \lambda(\mathcal{H}_{t-1,v}) \leq \frac{c_{t-1,v}}{(t-1)!}$ for all i. Thus

$$\lambda(\mathcal{H}_{t,v}) \leq \frac{c_{t-1,v}}{t!} \sum_{i=1}^{K} x_i (1-x_i)^{t-1}.$$

Since the support must span at least one edge of $\mathcal{H}_{t,v}$, we have $K \geq t \geq 2$. Moreover, since x is a legal weighting, $x_i \geq 0$ for all i, and $\sum_i x_i = 1$. Hence we may apply Lemma 2.8 to deduce that $\lambda(\mathcal{H}_{t,v}) \leq \frac{c_{t-1,v}}{t!}\left(1 - \frac{1}{K}\right)^{t-1}$. As $K \leq \frac{v^t-1}{v-1}$, this can be further bounded by $\lambda(\mathcal{H}_{t,v}) \leq \frac{c_{t-1,v}}{t!}\left(1 - \frac{v-1}{v^t-1}\right)^{t-1}$. Substituting in the definition of $c_{t-1,v}$, we obtain

$$\begin{aligned}
\lambda(\mathcal{H}_{t,v}) &\leq \frac{c_{t-1,v}}{t!}\left(1 - \frac{v-1}{v^t-1}\right)^{t-1} = \frac{1}{t!}\left(\prod_{i=0}^{t-2} \frac{v^{t-1}-v^i}{v^{t-1}-1}\right)\left(1 - \frac{v-1}{v^t-1}\right)^{t-1} \\
&= \frac{1}{t!}\left(\prod_{i=0}^{t-2} \frac{v^{t-1}-v^i}{v^{t-1}-1}\right)\left(\frac{v^t-v}{v^t-1}\right)^{t-1} = \frac{1}{t!}\prod_{i=0}^{t-2} \frac{\left(v^{t-1}-v^i\right)\left(v^t-v\right)}{\left(v^{t-1}-1\right)\left(v^t-1\right)} \\
&= \frac{1}{t!}\prod_{i=0}^{t-2} \frac{v^t-v^{i+1}}{v^t-1} = \frac{1}{t!}\prod_{i=1}^{t-1} \frac{v^t-v^i}{v^t-1} = \frac{1}{t!}\prod_{i=0}^{t-1} \frac{v^t-v^i}{v^t-1} = \frac{c_{t,v}}{t!},
\end{aligned}$$

completing the proof. □

3 An optimal construction

In this section we prove the lower bounds of Theorem 1.7 by providing, for every prime power v, a linear algebraic construction of an array of size $v^t \times k$ that covers a large number of t-sets.

We begin by handling the case $k = \frac{v^t-1}{v-1}$. Let $\mathbb{F}_v$ be the v-element field, and consider the vector space $\mathbb{F}_v^t$. Note that there are exactly $\frac{v^t-1}{v-1}$ 1-dimensional subspaces $L_1, \ldots, L_k$, and for each such subspace L_i, fix some non-zero vector $\vec{z}_i \in L_i \leq \mathbb{F}_v^t$. Now let A_{opt} be a $v^t \times \frac{v^t-1}{v-1}$ array whose rows are indexed by the v^t vectors in $\mathbb{F}_v^t$ and whose columns are indexed by $[k]$. Given $\vec{y} \in \mathbb{F}_v^t$ and $i \in [k]$, we define the $(\vec{y}, i)$ entry of A_{opt} to be the scalar product $\vec{y} \cdot \vec{z}_i$.

We claim that a t-set $Q \subseteq [k]$ is covered by A_{opt} if and only if the corresponding vectors $\{\vec{z}_i : i \in Q\}$ are linearly independent. To see this, let M be the $t \times t$ matrix containing the vectors $\{\vec{z}_i : i \in Q\}$ as columns. For any $\vec{y} \in \mathbb{F}_v^t$, the corresponding row of $(A_{\text{opt}})_Q$ is $\vec{y}^T M$. Now the set Q is covered by A_{opt} if every possible vector occurs in this way, or, equivalently, if any vector from $\mathbb{F}_v^t$ can be obtained as $\vec{y}^T M$ for some appropriate vector $\vec{y} \in \mathbb{F}_v^t$. This happens precisely when the matrix M is invertible, i.e. when the column vectors are linearly independent.

Now how many linearly independent sets $\{\vec{z}_{j_1}, \ldots, \vec{z}_{j_t}\}$ are there? If we have already chosen $i \geq 0$ linearly independent vectors $\vec{z}_{j_1}, \ldots, \vec{z}_{j_i}$, the next vector $\vec{z}_{j_{i+1}}$, and hence the corresponding subspace $L_{j_{i+1}}$ it belongs to, cannot be in the i-dimensional subspace spanned

by $\{\vec{z}_{j_1}, \ldots, \vec{z}_{j_i}\}$. This forbids $\frac{v^i-1}{v-1}$ of the possible vectors, leaving $\frac{v^t-1}{v-1} - \frac{v^i-1}{v-1}$ choices for $\vec{z}_{j_{i+1}}$. As the order in which the vectors are chosen does not matter, this gives a total of

$$\frac{1}{t!}\prod_{i=0}^{t-1}\left(\frac{v^t-1}{v-1} - \frac{v^i-1}{v-1}\right) = \frac{1}{t!}\prod_{i=0}^{t-1}\frac{v^t-v^i}{v-1} = \left(\frac{v^t-1}{v-1}\right)^t \frac{c_{t,v}}{t!} = c_{t,v}\frac{k^t}{t!}$$

different linearly independent sets of size t, and hence this is also the number of t-sets covered by A_{opt}. Note that this exactly matches the upper bound from Section 2, which implies that in Proposition 2.5, we in fact determine the Lagrangian of the hypergraph $\mathcal{H}_{t,v}$ precisely for all prime powers v.

If k is divisible by $\frac{v^t-1}{v-1}$, then we can take a blow-up of this linear algebraic construction, similar to the block construction in Section 1. Partition the set $[k]$ of column indices into $\frac{v^t-1}{v-1}$ parts of equal size, so that we have parts $\left\{B_i : i \in \left[\frac{v^t-1}{v-1}\right]\right\}$ with $|B_i| = \frac{k(v-1)}{v^t-1}$ for all i. We can now define the blown-up array $A_{\text{opt,block}}$ of size $v^t \times k$, where for $j \in B_i$, the jth column of $A_{\text{opt,block}}$ is the ith column of A_{opt}.

It is easy to see that a set $Q \subseteq [k]$ is covered by $A_{\text{opt,block}}$ if and only if it contains at most one element from each block and the corresponding set of block indices is covered by A_{opt}. Thus

$$\text{Cov}(A_{\text{opt,block}}) = \bigcup_{Q' \in \text{Cov}(A_{\text{opt}})} \prod_{i \in Q'} B_i.$$

Restricting to sets of size t, we find

$$\begin{aligned}
\text{cov}_t(A_{\text{opt,block}}) &= \sum_{\{j_1,\ldots,j_t\} \in \text{Cov}(A_{\text{opt}})} \prod_{i=1}^{t} |B_{j_i}| \\
&= \sum_{\{j_1,\ldots,j_t\} \in \text{Cov}(A_{\text{opt}})} \left(\frac{k(v-1)}{v^t-1}\right)^t = \left(\frac{k(v-1)}{v^t-1}\right)^t \text{cov}_t(A_{\text{opt}}) \\
&= \left(\frac{k(v-1)}{v^t-1}\right)^t \left(\frac{v^t-1}{v-1}\right)^t \frac{c_{t,v}}{t!} = c_{t,v}\frac{k^t}{t!},
\end{aligned}$$

showing we do have equality in Theorem 1.7 for this case.

For general k, as noted by Alon (personal communication), one can take a random partition of the set $[k]$ of column indices into the $\frac{v^t-1}{v-1}$ parts $\left\{B_i : i \in \left[\frac{v^t-1}{v-1}\right]\right\}$. More precisely, for each $j \in [k]$ choose some $i \in \left[\frac{v^t-1}{v-1}\right]$ independently and uniformly at random, and add j to B_i. We can now take the corresponding block construction A based on A_{opt}, where for each $j \in B_i$, the jth column of A is the ith column of A_{opt}. We again have that a set Q is covered by A if and only if it contains at most one element from each block and the corresponding set of block indices is covered by A_{opt}. Accordingly, for a fixed t-set Q,

$$\mathbb{P}(Q \in \text{Cov}(A)) = \frac{t!\,\text{cov}_t(A_{\text{opt}})}{\left(\frac{v^t-1}{v-1}\right)^t} = c_{t,v},$$

and so by linearity of expectation the expected number of t-sets covered by A is $c_{t,v}\binom{k}{t}$. Hence there must be some block partition for which the number of covered t-sets is at least $c_{t,v}\binom{k}{t}$. This concludes the proof of Theorem 1.7.

To close this section, we note that in the special case $v = 2$, the columns appearing in A_{opt} are, up to an affine translation, the non-constant columns of the $2^t \times 2^t$ Hadamard matrix constructed by Sylvester [30]. While this may not be immediately apparent from the recursive definition of these matrices, it follows from an equivalent formulation given in [4]. Here, the rows and columns of the matrix are indexed by vectors from $\mathbb{F}_2^t$, and the entry in the row corresponding to $\vec{y}$ and the column corresponding to $\vec{z}$ is $(-1)^{\vec{y}\cdot\vec{z}}$.

4 Small covering arrays

In this section we shall build small covering arrays, thereby proving Corollary 1.8. Roughly speaking, the idea is to combine many random copies of the construction from Section 3 into a single array. Since a t-set Q is covered with large probability by each individual copy of the small construction, it follows that the probability that Q is not covered by any of the copies will be exponentially small. In order to show that we obtain a covering array, we need to prove that with positive probability, all of the $\binom{k}{t}$ t-sets of columns are covered. This will follow from the Lovász Local Lemma, first proven by Erdős and Lovász [6] and subsequently sharpened by Spencer [29].

Theorem 4.1 (Lovász Local Lemma, 1975)**.** *Let $\mathcal{E}_1, \mathcal{E}_2, \ldots, \mathcal{E}_m$ be events in some probability space. Suppose there are $p \in [0,1]$ and $d \in \mathbb{N}$ such that for each $i \in [m]$, $\mathbb{P}(\mathcal{E}_i) \leq p$, and the event $\mathcal{E}_i$ is mutually independent of a set of all but at most d of the other events. If $ep(d+1) \leq 1$, then $\mathbb{P}\left(\cap_{i=1}^m \overline{\mathcal{E}_i}\right) > 0$.*

We can now proceed with the proof of Corollary 1.8.

Proof of Corollary 1.8. We first handle the prime power case. Suppose we have integers $k \geq t \geq 2$, and let $v \geq 2$ be a prime power. We wish to build an $(rv^t; t, k, v)$-covering array A, where $r = \left\lceil \frac{(t-1)\log_2 k + \log_2(et)}{\log_2 \frac{1}{1-c_{t,v}}} \right\rceil$.

For $1 \leq \ell \leq r$, let $A^{(\ell)}$ be an independent copy of our construction from Section 3. That is, for each $j \in [k]$, let $i_{\ell,j} \in \left[\frac{v^t-1}{v-1}\right]$ be chosen independently and uniformly at random, and take the jth column of $A^{(\ell)}$ to be $(A_{\text{opt}})_{\{i_{\ell,j}\}}$. We then take A to be the concatenation of $A^{(1)}, A^{(2)}, \ldots, A^{(r)}$, giving an $rv^t \times k$ array. We wish to show that with positive probability A is an $(rv^t; t, k, v)$-covering array.

Hence, for each t-subset Q of the columns, we let $\mathcal{E}_Q$ be the event that Q is *not* covered by A. Since the array A contains each array $A^{(\ell)}$, $\ell \in [r]$, if Q is not covered by A, then it is not covered by any $A^{(\ell)}$. From Section 3, we have $\mathbb{P}\left(Q \notin \text{Cov}(A^{(\ell)})\right) = 1 - c_{t,v}$, and by construction these subarrays are independent of one another. We may therefore define p by

$$\mathbb{P}(\mathcal{E}_Q) = \mathbb{P}(Q \notin \text{Cov}(A)) \leq \mathbb{P}\left(\cap_{\ell\in[r]} \left\{Q \notin \text{Cov}\left(A^{(\ell)}\right)\right\}\right) = (1 - c_{t,v})^r = p.$$

Our construction of the array A also ensures that the different columns are independent of one another. In particular, this implies that the event $\mathcal{E}_Q$ is mutually independent of the set of all events that depend on a disjoint set of columns, i.e. $\{\mathcal{E}_{Q'} : Q' \cap Q = \emptyset\}$. For a fixed set Q, if Q' intersects Q, it contains one of the t columns of Q, and then there are fewer than $\binom{k}{t-1}$ choices for the remaining columns of Q'. Hence we may take $d < t\binom{k}{t-1}$.

We thus have

$$ep(d+1) \leq et\binom{k}{t-1}(1-c_{t,v})^r \leq etk^{t-1}(1-c_{t,v})^r \leq 1,$$

where the final inequality follows from our choice of r. Hence, by Theorem 4.1, with positive probability none of the events $\mathcal{E}_Q$ occur, which implies the existence of some such $rv^t \times k$ array A that covers all of its $\binom{k}{t}$ t-subsets of columns. This gives the claimed bound,

$$\mathrm{CAN}(t,k,v) \leq rv^t = \left\lceil \frac{(t-1)\log_2 k + \log_2(et)}{\log_2 \frac{1}{1-c_{t,v}}} \right\rceil v^t.$$

For the next bound, we divide the above expression by $\log_2 k$ and take the limit as k tends to infinity, giving

$$d(t,v) \leq \frac{(t-1)v^t}{\log_2 \frac{1}{1-c_{t,v}}} \leq \frac{(t-1)v^t}{2\log_2 v - \log_2(v+1)},$$

where the second inequality follows from the bound $c_{t,v} \geq 1 - \frac{v+1}{v^2}$, derived below the statement of Theorem 1.7.

Finally, we turn to the case when v is not a prime power. Here we use the trivial observation that for any $v \leq v'$, $\mathrm{CAN}(t,k,v) \leq \mathrm{CAN}(t,k,v')$, since projecting an array from a large set of symbols to a smaller set cannot cause any subset of columns to become uncovered. In particular, it follows that $d(t,v) \leq d(t,v')$.

Given v, let q be the smallest prime power that is at least v. Baker, Harman and Pintz [2] proved that, provided v is sufficiently large, $v \leq q \leq v + v^{0.526}$. We thus have

$$d(t,v) \leq d(t,q) \leq \frac{(t-1)q^t}{2\log_2 q - \log_2(q+1)} \leq \frac{(t-1)(1+v^{-0.474})^t v^t}{2\log_2 v - \log_2(v+1)} \leq \frac{(t-1)e^{tv^{-0.474}}v^t}{2\log_2 v - \log_2(v+1)},$$

as required. □

One may contrast our methods with those that have been used in previous constructions. In [27] and [28], small covering arrays were built by algebraically extending a random array. Here, we do the opposite, taking random copies of a small linear algebraic array. Since our initial array was exceedingly efficient in covering t-sets, we were able to reduce the number of random copies required, thus resulting in a smaller construction. Indeed, the upper bound of Section 2 shows that this is the best $v^t \times k$ array one can start with. However, to further improve the upper bound, one could perhaps find better ways to combine these copies, or perhaps start with a larger structured construction.

5 Concluding remarks

In this paper we combined linear algebraic and probabilistic arguments to construct small covering arrays, asymptotically improving the upper bounds on $\mathrm{CAN}(t,k,v)$ by a factor of $\ln v$. This involved the study of the extremal function $\mathrm{cov}_{\max}(N;t,k,v)$, and we showed that at the lower threshold $N=v^t$ (the minimum size of an array that permits covering a single t-set) one can already cover a large proportion of all t-sets. We close with some final remarks and possible directions for further research.

Almost-covering arrays As mentioned in the introduction, Sarkar, Colbourn, De Bonis and Vaccaro [28] studied (in greater generality) almost-covering arrays. Given N rows, k columns, a strength t, a set S of v symbols, and a coverage fraction $\varepsilon \in [0,1]$, an array $A \in S^{N\times k}$ is an $(N;t,k,v,\varepsilon)$-almost-covering array if A covers all but at most $\varepsilon\binom{k}{t}$ t-subsets Q. This relaxes the concept of a covering array, and one can again define an extremal function $\mathrm{ACAN}(t,k,v,\varepsilon)$, which is the smallest N for which an $(N;t,k,v,\varepsilon)$-almost-covering array exists.

This is the inverse function of the $\mathrm{cov}_{\max}(N;t,k,v)$ function; rather than fixing the size of the array and maximising the number of t-sets covered, we fix the number of sets to be covered, and minimise the size of the array needed. In [28] it is proven that almost-covering arrays can be significantly smaller than covering arrays; indeed, the number of rows need not grow with k. More precisely, they showed $\mathrm{ACAN}(t,k,v,\varepsilon) \le v^t \ln\left(\frac{v^{t-1}}{\varepsilon}\right)$, and that when v is a prime power, the bound can be improved to $\mathrm{ACAN}(t,k,v,\varepsilon) \le v^t \ln\left(\frac{2v^{t-2}}{\varepsilon}\right) + v$.

As in Section 4, we can improve the bounds by a $\log_2 v$ factor by concatenating random copies of the linear algebraic array A_{opt} until the expected number of uncovered sets is at most $\varepsilon\binom{k}{t}$. For prime power v, this gives $\mathrm{ACAN}(t,k,v,\varepsilon) \le v^t \left\lceil \frac{\ln \frac{1}{\varepsilon}}{2\ln v - \ln(v+1)} \right\rceil$. In particular, A_{opt} itself shows $\mathrm{ACAN}(t,k,v,\varepsilon) = v^t$ for $\varepsilon \ge \frac{v+1}{v^2}$. These results can again be extended to all large values of v by replacing v with the next prime power, and this gives effective bounds whenever $t = o(v^{0.474})$.

Explicit constructions of covering arrays To improve the bound on $\mathrm{CAN}(t,k,v)$, we consider a random construction and apply the Lovász Local Lemma, making our result an existential one. However, we hope that one could apply the algorithmic version of the Local Lemma, due to Moser and Tardos [19], in our setting, which would perhaps result in an efficient Las Vegas algorithm to produce small covering arrays. For a simpler analysis, one could take a slightly larger number of random copies of A_{opt}, which would then give a covering array with high probability, resulting in an efficient Monte Carlo algorithm instead.

Regardless, in light of the many important applications of covering arrays, a great deal of interest lies in the explicit construction of efficient covering arrays. In the binary setting, with $v=2$, where randomised constructions give arrays of size $\mathcal{O}\left(t2^t \log_2 k\right)$, Alon [1] gave an explicit construction of size $2^{\mathcal{O}(t^4)} \log_2 k$. Following a sequence of incremental improvements, Naor, Schulman and Srinivasan [21] provided near-optimal arrays of size $t^{\mathcal{O}(\log_2 t)} 2^t \log_2 k$.

Given the inherent symmetry of the linear algebraic array A_{opt}, it might be possible to deterministically concatenate copies of this array to form a small explicitly-constructed covering array. Indeed, this might even give better upper bounds on $\text{CAN}(t, k, v)$ than those we obtained in Section 4 using probabilistic means. Taking this approach, the minimum number of copies of A_{opt} needed is given by the following hashing problem with an algebraic twist.

Question 5.1. *Given $k \geq t \geq 2$ and a prime power $v \geq 2$, what is the minimum r such that there is a collection of k sequences in $(\mathbb{F}_v^t)^r$ (that is, the sequences have length r, and each entry is a vector in $\mathbb{F}_v^t$) with the property that for every t-subset of the sequences, there is some coordinate where the t corresponding vectors are linearly independent?*

In the standard hashing problem, one would only require the existence of a coordinate where the sequences were pairwise-distinct, whereas here we impose a stronger algebraic condition. Given that the best-known bounds for the hashing problem come from probabilistic constructions (see [9]), it might be difficult to improve on the bounds from Section 4 this way. However, deterministic solutions to Question 5.1 would result in explicit constructions of (hopefully) small covering arrays.

Evolution of $\text{cov}_{\max}(N; t, k, v)$ Another way to improve the construction in Section 4 would be to replace A_{opt}, and instead concatenate random copies of some other array. The upper bound proven in Section 2 shows that we cannot hope to do better with arrays of size v^t, but it might be beneficial to consider initial arrays with a larger number of rows. With this in mind, it would be of interest to determine how the function $\text{cov}_{\max}(N; t, k, v)$ grows as N increases. In particular, at the other extreme we observe that $\text{CAN}(t, k, v) = \min\left\{N : \text{cov}_{\max}(N; t, k, v) = \binom{k}{t}\right\}$, and so complete knowledge of $\text{cov}_{\max}(N; t, k, v)$ would solve the covering array problem as well.

On a much finer scale, what happens for $N = v^t + s$ for small values of s? In A_{opt}, if a t-subset Q of columns is not covered, then the vectors $\vec{z}_i \in \mathbb{F}_v^t$ that the columns are mapped to form a matrix of rank at most $t-1$, and hence at most v^{t-1} rows appear in the subarray $(A_{\text{opt}})_Q$. This implies that to increase the number of covered t-subsets, we need to add at least $v^t - v^{t-1}$ new rows, almost doubling the size of A_{opt}.

Of course, there could be other arrays of size $v^t + s$ that do not contain A_{opt} as a subarray, but cover a larger number of t-sets. Our proof of the optimality of A_{opt} when $N = v^t$ relied heavily on the fact that if a t-set is covered, then each sequence in S^t appears exactly once as a row in the corresponding subarray, leading to useful linear algebraic interpretations of coverage. This rigid structure is lost for larger arrays, rendering the analysis more difficult. Still, we feel that without this structure, one should not be able to cover a larger number of t-sets. To make this intuition more precise, we offer the following question.

Question 5.2. *When v is a prime power, are $\text{cov}_{\max}(v^t + 1; t, k, v)$ and $\text{cov}_{\max}(v^t; t, k, v)$ equal? What is the smallest s for which $\text{cov}_{\max}(v^t + s; t, k, v) > \text{cov}_{\max}(v^t; t, k, v)$?*

Non-uniform coverage Finally, when dealing with covering arrays, we have only focussed on the number of subsets of some fixed size t that are covered. However, one could instead consider all covered sets, and look to maximise $|\mathrm{Cov}(A)|$ instead.

In the context of set shattering, similar questions have indeed been considered. The famous Sauer–Shelah inequality states that a set family $\mathcal{F}$ must shatter at least $|\mathcal{F}|$ sets in total. Given its numerous applications, a pressing open problem is the classification of all families that attain this bound with equality. For details on this line of research, see, for example, [3, 18, 26].

For our problem, we instead ask which families of a given size maximise the number of shattered sets. Note that a family of size m can shatter sets of size at most $\lfloor \log_2 m \rfloor$, and, if $\log_2 m$ is small compared to the size k of the ground set, then almost all such sets will have size precisely $\lfloor \log_2 m \rfloor$. If $m = 2^t$, our construction maximises the number of t-sets shattered, and hence one might expect it also maximises the total number of shattered sets.

Question 5.3. *Given $0 \leq m \leq 2^k$, which set family $\mathcal{F} \subseteq 2^{[k]}$ of size $|\mathcal{F}| = m$ maximises the number of shattered sets?*

Acknowledgements The authors would like to express their gratitude towards Noga Alon, Gal Kronenberg and Benny Sudakov for helpful discussions at various stages of this project. Special thanks go to Ugo Vaccaro for bringing [28] to our attention, prompting the extension of our initial results on set shattering to general covering arrays.

References

[1] N. Alon, *Explicit construcion of exponential sized families of k-independent sets*, Discrete Math. **58.2** (1986), 191–193.

[2] R. C. Baker, G. Harman and J. Pintz, *The difference between consecutive primes, II*, P. Lond. Math. Soc. **83.3** (2001), 532–562.

[3] B. Bollobás and A. J. Radcliffe, *Defect Sauer results*, J. Comb. Theory A **72** (1995), 189–208.

[4] P. J. Cameron (ed.), *Hadamard matrices*, Encyclopaedia of Design Theory, `http://www.maths.qmul.ac.uk/~leonard/designtheory.org/library/encyc/topics/had.pdf`.

[5] C. J. Colbourn, *Combinatorial aspects of covering arrays*, Le Matematiche (Catania) **58** (2004), 121–167.

[6] P. Erdős and L. Lovász, *Problems and results on 3-chromatic hypergraphs and some related questions*, Infinite and Finite Sets (A. Hajnal, R. Rado and V. T. Sós, eds.), North-Holland, Amsterdam (1975), 609–628.

[7] N. Francetić and B. Stevens, *Asymptotic size of covering arrays: an application of entropy compression*, J. Comb. Des. (2016), doi: 10.1002/jcd.21553.

[8] P. Frankl and V. Rödl, *Hypergraphs do not jump*, Combinatorica **4** (1984), 149–159.

[9] M. L. Fredman and J. Komlós, *On the size of separating systems and families of perfect hash functions*, SIAM J. Algebra. Discr. **5.1** (1984), 61–68.

[10] L. Gargano, J. Körner and U. Vaccaro, *Sperner capacities*, Graph. Combinator. **9** (1993), 31–46.

[11] A. P. Godbole, D. E. Skipper and R. A. Sunley, *t-covering arrays: upper bounds and Poisson approximations*, Comb. Probab. Comput. **5** (1996), 105–118.

[12] A. Hartman and L. Raskin, *Problems and algorithms for covering arrays*, Disc. Math. **284** (2004), 149–156.

[13] A. S. Hedayat, N. J. A. Sloane and J. Stufken, *Orthogonal Arrays: Theory and Applications*, Springer Science and Business Media (1999).

[14] Gy. O. H. Katona, *Two applications (for search theory and truth functions) of Sperner type theorems*, Period. Math. Hungarica **3.1-2** (1973), 19–26.

[15] D. J. Kleitman and J. Spencer, *Families of k-independent sets*, Discrete Math. **6.3** (1973), 255-262.

[16] J. Lawrence, R. N. Kacker, Y. Lei, D. R. Kuhn and M. Forbes, *A survey of binary covering arrays*, Electron. J. Combin. **18.1** (2011), P84.

[17] J. R. Maximoff, M. D. Trela, D. R. Kuhn and R. Kacker, *A method for analyzing system state-space coverage within a t-wise testing framework*, 4th Annual IEEE Systems Conference (2010), 598–603.

[18] T. Mészáros and L. Rónyai, *Shattering-extremal set systems of VC dimension at most 2*, Electron. J. Combin. **21.4** (2014) P4.30.

[19] R. A. Moser and G. Tardos, *A constructive proof of the general Lovász Local Lemma*, J. ACM **57.2** (2010), 11.

[20] T. S. Motzkin and E. G. Straus, *Maxima for graphs and a new proof of a theorem of Turán*, Canad. J. Math. **17** (1965), 533–540.

[21] M. Naor, L. J. Schulman and A. Srinivasan, *Splitters and near-optimal derandomization*, Foundations of Computer Science, Proceedings, IEEE (1995), 182–191.

[22] R. L. Plackett and J. P. Burman, *The design of optimum multifactorial experiments*, Biometrika **33.4** (1946), 305–325.

[23] C. R. Rao, *Hypercubes of strength d leading to confounded designs in factorial experiments*, Bull. Calcutta Math. Soc. **38** (1946), 67–78.

[24] C. R. Rao, *Factorial experiments derivable from combinatorial arrangements of arrays*, J. Roy. Stat. Soc., Suppl. **9** (1947), 128–139.

[25] A. Rényi, *Foundations of Probability*, Wiley, New York (1971).

[26] L. Rónyai and T. Mészáros, *Some combinatorial applications of Gröbner bases*, International Conference on Algebraic Informatics, Springer Berlin Heidelberg (2011), 65–83.

[27] K. Sarkar and C. J. Colbourn, *Upper bounds on the size of covering arrays*, arXiv:1603.07809 (2016).

[28] K. Sarkar, C. J. Colbourn, A. De Bonis and U. Vaccaro, *Partial covering arrays: algorithms and asymptotics*, International Workshop on Combinatorial Algorithms, Springer International Publishing (2016), 437–448.

[29] J. H. Spencer, *Asymptotic lower bounds for Ramsey functions*, Discrete Math. **20** (1977), 69–76.

[30] J. J. Sylvester, *Thoughts on inverse orthogonal matrices, simultaneous sign successions, and tessellated pavements in two or more colours, with applications to Newton's rule, ornamental tile-work, and the theory of numbers*, Philos. Mag. **34** (1867), 461–475.